Persistence Mechanics: Architecture, Dynamics, and the Entropic Action

Kate Lenore Meyer
*Independent Researcher**
(Dated: September 3, 2026)

Biological organisms, ecological networks, engineered feedback loops, and computational processes all persist by solving the same problem: maintaining structural coherence against entropic erasure. We formalize the mechanics of persistence, deriving the substrate-independent, minimal architecture of persistence at the zero flux limit from information-theoretic first principles. We show that any persistent system must solve three existential problems: Finiteness (bounded resources), Conservation (preserved information), and Causality (well-defined evolution), yielding six irreducible pillars: Capacity, Map, Protocol, Governor, Toll, and Margin. We ground this static architecture in time by deriving the minimum dynamical description from three irreducible primitives of state transition: irreducible toll (c), structural complexity (K), and coordinated transition flux (f). Their product is the dissipation density $\mathcal{D} = c \cdot K \cdot f$. We derive the Entropic Action S_Φ, the cumulative integral of $\mathcal{D}$, and demonstrate that the Principle of Least Action emerges as the boundary of survivorship: non-minimizing trajectories undergo exponential suppression, leaving only the least-action remainder as observable history. Mapping diverse systems onto this architectural invariant recasts domain-specific heuristics as instances of the same universal mechanics: solutions validated in one domain transfer to another, persistent systems become directly comparable, and structural failure becomes a quantifiable, cross-domain phenomenon.

I. INTRODUCTION

From biological organisms to engineered feedback loops, complex systems science describes how structures maintain coherence against entropic disturbances [1]. Their defining feature is *persistence*; they retain their structure despite ongoing noise and decay. Control theory describes how systems regulate against fluctuations, taking the regulation target (setpoint) as an externally specified boundary condition. Information theory quantifies the cost of information processing, taking the processor as given. Thermodynamics establishes the arrow of time, treating entropy as a property of states. These three fields developed independently to solve distinct engineering problems: thermodynamics from heat engines, information theory from telecommunication, and control theory from servomechanisms. Each is optimizing a system that already exists.

We define this missing foundation: the minimal architecture presupposed by any dynamical or optimization discipline. Grounding this investigation in algorithmic information theory, we identify the invariant substrate that information theory; thermodynamics and control theory each presuppose. By uniting this static architecture with the minimal dynamic action required to maintain it, we establish *persistence mechanics*: the necessary conditions for a single persistent system, its dynamics, and the selection principle governing observable history.

Previous complex systems research has converged on pieces of this puzzle from multiple directions. Prigogine established that persistent structures occupy a minimum-entropy-production attractor[2]; Friston formulated persistence as the minimization of variational free energy via an internal model [3]; England showed how external drive selects for structural configurations in non-equilibrium steady states[4]; and cybernetics has catalogued the organizational patterns, such as requisite variety and good regulation, that correlate with survival[5, 6]. These contributions provide convergent evidence for the components of persistence, yet each treats the structure to be persisted as given. The minimal conditions required for a persistent structure to exist remain unformalized.

We describe two universal mechanics found in systems that persist. First, the **Architecture of Persistence** (Section II): six substrate-independent pillars derived from the logical requirements of Finiteness, Conservation, and Causality. This architecture specifies what must exist, not how it behaves under transition. Second, the **Dynamics of Persistence** (Section III): the universal behavior governing how any persistent system dissipates, accumulates entropic action, and undergoes selection. We derive the dissipation density $\mathcal{D} = c \cdot \mathbf{K} \cdot f$ as the minimum dynamical description any state transition law must presuppose.

Dynamics determines which histories remain observable. Section IV shows that the Principle of Least Action emerges as the boundary of survivorship: trajectories that do not minimize their Entropic Action undergo exponential suppression, leaving only the least-action remainder as observable history. Section V extends the architecture to open systems and section VI then illustrates the architecture's mapping across biological, ecological, and computational substrates.

II. THE ARCHITECTURE OF PERSISTENCE

Any system that maintains structural coherence against entropic decay must solve three existential prob-

* kate@meyerhome.net

lems. The modes of engagement with these problems yield six irreducible pillars, which we derive, quantify, and prove minimally necessary and sufficient by logical exhaustive partitioning.

A. The Structural Requirements

Any process persisting through time is bounded in extent, carries distinguishable content, and evolves through an ordered sequence of states. Absent any one, it is indistinguishable from its environment, from noise, or from something static. Formally, these correspond to: **Finiteness** (bounded resources against divergence), **Conservation** (preserved information against dissipation), and **Causality** (well-defined temporal evolution against ambiguity).

Existence, however, is transient. To persist, each requirement must survive its characteristic perturbation.

Each therefore demands **provision** (establishing the means to satisfy the requirement) and **defense** (protecting that means from the characteristic failure mode) yielding six irreducible pillars.

Requirement	Provision	Defense
Finiteness	Capacity	Margin
Conservation	Map	Protocol
Causality	Toll	Governor

Table I. By defining a persisting system as an entity possessing finite extent, conserved content, and causal sequence, and defining survival as the provisioning and defense of these properties, we establish a logically closed $3 \cdot 2$ taxonomy. This yields six necessary pillars.

These pillars operate upon a medium that imposes immutable limits (bandwidth, latency, noise floor) known as the **Substrate**.

B. The Six Pillars

The previous section argued what must exist by logical exhaustion. This section shows that each pillar has a convergent theorem from six unrelated fields.

Capacity (CAP): *The Vessel.* The infrastructure required to acquire resources, process signals, and perform work. It defines the system's operational potential. *Convergent Evidence (The Bekenstein Bound [7]):* Bekenstein's bound limits the information content of any finite region, and the Margolus-Levitin limit[8] bounds the rate of state transitions; together they establish that finite capacity is a structural invariant.

Margin (MAR): *The Buffer.* The structural and energetic reserves held to absorb stochastic shocks. It defines the system's 'Safe Operating Area,' providing the buffer required to survive environmental variance without terminal failure. (The gain and phase margin). *Convergent Evidence (Nyquist-Bode Stability Margins[9] and Kingman's Formula[10]):* A system operating at theoretical criticality has zero tolerance for variance. Kingman's formula shows that as a queueing system's utilization approaches unity, waiting time diverges asymptotically, zero buffer against arrival variance means unbounded delay. Nyquist's stability margins are the linear-control manifestation of the same general requirement: persistent systems must maintain structural tolerance away from critical thresholds.

Map (MAP): *The Model.* The internal logic, topological structure, or compressed representation of environmental regularities that distinguishes the system from the environment. It functions as a predictive engine that reduces uncertainty (informational surprise), allowing the system to steer towards coherent states. *Convergent Evidence (The Good-Regulator Theorem [6]):* Conant and Ashby's Good-Regulator theorem establishes that persistent regulation requires an internal model; without Map, the system–environment boundary dissolves.

Protocol (PRO): *The Interface.* The communicative glue regulating information flow between the Capacity and the Map. It ensures internal coherence, standardizes interfaces, and minimizes signal decay across the system's internal boundaries. *Convergent Evidence (Shannon's Source-Channel Separation and Channel Capacity [11]):* Shannon's source–channel separation theorem shows that coherence across internal boundaries requires an independent interface layer; Protocol is that layer.

Toll (TOL): *The Temporal Cost.* State transitions carry an irreducible information and temporal cost. Any irreversible logical operation, such as erasing a bit of structural information, necessitates a strictly compensatory increase in environmental Shannon entropy (the $k_B T \ln 2$ bound). *Convergent Evidence (Landauer's Principle and Margolus-Levitin Limit [8, 12]):* These bounds make Toll a structural requirement.

Governor (GOV): *The Constraint.* The non-linear constraint mechanism that prevents unbounded divergence or 'runaway' loops. It enforces operational boundaries (e.g., apoptosis, budget limits) to protect the substrate from exhaustion. *Convergent Evidence (Ashby's Law of Requisite Variety [5]):* Ashby's Law of Requisite Variety requires that the constraint mechanism match the variety of perturbations; without Governor, runaway is unavoidable.

With each pillar independently attested, the architecture becomes grounded with citable bounds and can be built on formally.

C. Necessity: Counterfactual Failure Modes

The necessity of the set $\mathbb{P} = \{\text{CAP, MAP, PRO, GOV, TOL, MAR}\}$ follows by analyzing the **Counterfactual Failure Mode** of a system lacking exactly one pillar ($\mathbb{P}' = \mathbb{P} \setminus \{x\}$). In every case, the expected lifetime of the system τ tends to zero.

- **No Capacity (Starvation):** Without baseline operational infrastructure, the system cannot acquire, process, or export resources. Once existing reserves are spent, any perturbation is fatal; $\tau \to 0$ via rapid exhaustion.

- **No Margin (Fragility):** Operating at theoretical criticality leaves zero tolerance for environmental variance. Any non-zero perturbation exceeds the system's structural limits immediately, so the expected time to a fatal shock is finite; $\tau \to 0$ by random shock.

- **No Map (Dissolution):** Without an internal model, the system has no reference for coherence, and no boundary to maintain against its environment; $\tau \to 0$ as the system–environment boundary dissolves.

- **No Protocol (Decoherence):** Components act on statistically independent, outdated states, with no mechanism to reconcile them; $\tau \to 0$ by organizational fragmentation.

- **No Toll (Zeno Collapse):** A causal sequence requires irreducibly discrete state transitions. Without an irreducible toll per operation, the sequence can be infinitely subdivided; no individual step remains distinguishable as a cause of the next, and the ordering required for well-defined evolution dissolves. $\tau \to 0$ by causal decoherence.

- **No Governor (Divergence):** Unconstrained positive feedback drives resource consumption toward unbounded growth, exhausting the substrate or diverging past any recoverable bound; $\tau \to 0$ by unbounded runaway.

Since the removal of any component results in $\tau \to 0$, every pillar is necessary.

D. Sufficiency: Compositional Closure

Proving sufficiency over an unbounded architectural space would require ruling out every possible construction, which no finite argument can do. What we can show instead is that every capability traditionally treated as an independent requirement decomposes into the six pillars, with no counterexample identified to date. Memory is Map content provisioned against temporal variance by Margin. Prediction is Map operation on incoming signal via Protocol. Repair is Governor-triggered re-provisioning within the Margin reserve. Replication is Protocol-mediated write of Map onto fresh Capacity. Advanced topologies such as feedforward networks, state observers, and adaptive controllers extend the six-pillar structure the same way: none introduces a new primitive, each is a composition within the existing six. No counterexample or seventh independent primitive has been identified.

E. The Cost of Persistence

The six pillars describe what must be in place for a system to persist, but maintaining that architecture consumes resources: a persistent system is a resistor against entropic erosion, paying continuously just to remain distinguishable from its environment. To compare systems or define what minimizing means, we need that total structural cost as a single scalar.

Every pillar imposes a strictly positive cost: a system cannot maintain boundaries, store internal models, or execute state updates for free. Within this strictly positive budget, the pillars partition into two operational roles:

- **Direct Structural Costs:** The irreducible physical overhead of maintaining the system's baseline existence. Maintaining the vessel (Z_{CAP}, Z_{MAR}), the internal model (Z_{MAP}), and executing state updates (Z_{TOL}) requires a strict allocation of resources regardless of how the system is organized.

- **Coordination Costs:** Protocol and Governor incur direct structural costs to operate, but they function as active constraints rather than passive structures. Their operational effect is to suppress uncoordinated variance.

The system pays a continuous *coordination cost* (maintaining Protocol and Governor) as an insurance premium to avoid a catastrophic entropic penalty (fragmentation or runaway divergence). In control theory, this is analogous to negative feedback: the active controller consumes continuous energy to operate, but its regulation ensures the closed loop's total dissipation is lower than the open-loop alternative.

To quantify this baseline structural cost without conflating it with environmental work, we isolate the system at its zero-flux limit. In Quiescent Steady State (QSS), no net energy or information enters or leaves. The six pillars operate on functionally orthogonal domains, so their contributions do not interfere. The Net Entropic Impedance is therefore a strictly additive budget across these independent channels:

$$\bar{Z}_{\text{net}} = Z_{\text{CAP}} + Z_{\text{MAP}} + Z_{\text{TOL}} + Z_{\text{MAR}} + Z_{\text{PRO}} + Z_{\text{GOV}} \quad (1)$$

Because the pillars occupy orthogonal domains at zero flux, any correlation between them is already priced into the individual coordination terms Z_{PRO} and Z_{GOV}; no additional cross-term is needed to capture coupling between pillars within this baseline. Coupling outside QSS introduces non-linear contention, which we formalize in section V.

F. Relation to Existing Disciplines

Previous disciplines have encountered pieces of the six-pillar architecture already. In each case the components were taken as given, not recognized as instances of one invariant beyond the discipline's own domain. The universal architecture frames these components as convergent evidence for a single structural requirement: each domain encountered the same architecture from its own direction, and the present mechanics makes that convergence explicit.

1. Relation to Information Theory

Channel optimization repeatedly encounters parameters corresponding to the six pillars: channel capacity (Capacity), codebook topology (Map), handshaking and framing (Protocol), power constraints (Governor), processing overhead (Toll), and link margins (Margin). Classical information theory treats these six as externally specified boundary conditions and optimizes throughput conditioned on them, without asking why a physical system must instantiate all six simply to persist against entropic erasure.

2. Relation to Thermodynamics

Thermodynamics assumes an external thermal bath and continuous ensemble variables from the outset; it defines temperature and equilibrium against that bath without asking what structure must exist prior to it. Algorithmic information theory supplies the missing formal substrate. Shannon entropy and Gibbs entropy are formally identical in the statistical-mechanical limit, while Landauer's principle and the Margolus-Levitin bound supply the physical scale of individual state transitions: information loss corresponds to entropic dissipation, channel capacity to thermodynamic availability, and information flux to energy flux.

The difference is one of scope and substrate. The zero-flux architecture requires neither bath nor temperature. It is defined at zero flux on a discrete information substrate, and thermodynamic phenomenology emerges as the continuum limit of accumulated discrete transitions. Where thermodynamics describes how heat flows, this underlying architecture describes what structure must exist before heat or flow can be meaningfully defined.

3. Relation to Control Theory

Classical control theory emerged from engineering physical servomechanisms: radar mounts, autopilots, and chemical plants whose historical context was optimizing systems already built, powered, and staffed. Its four given components (the plant, the sensor, the controller, and the actuator) are convergent evidence for four of the six pillars: Plant/Actuator → Capacity, Reference Signal/Setpoint → Map, Sensor/Feedback Path → Protocol, Controller/Limiter → Governor.

The remaining two pillars are rarely canonized alongside the first four as they are not physical components, but describe how the four components behave under time and stress. Time delays, bandwidth limits, and processing latency (→ Toll) were minimized as engineering overhead and gain/phase margins (→ Margin) were tuned for stability.

Control theory asks how to regulate an existing system; the present mechanics derives the architecture that any existing system must instantiate. The classical four-component loop is not a rival account but one physical instantiation of the six-pillar invariant.

III. THE DYNAMICS OF PERSISTENCE

The static architecture of persistence necessitates a minimal dynamical description. A system that persists maintains structural coherence against entropic erosion; therefore it must export internal entropy to its environment at a rate that balances what it produces. We construct the measure of that dissipation from the three irreducible primitives that any dynamical law governing persistent systems must presuppose, prove their necessity, and integrate them into a cumulative Entropic Action that maps directly to the six-pillar architecture.

A. The Cost of State Transitions

While a system's architecture provides the capability to persist, the operation of those pillars determines how long it actually persists. Any state transition generates entropy; left unconstrained, a system's own activity can destabilize it.

As Prigogine established[2], persistence therefore requires two behaviors: internal entropy production must be exported to the environment fast enough that the system's own entropy does not rise without bound in the aggregate, and that production rate itself settles to the smallest value consistent with holding steady state ($\delta\sigma = 0$): below this floor, the system cannot maintain stability; above it, state transitions incur unnecessary dissipation.

We quantify this at its limiting case, QSS, where $\bar{Z}_{\mathrm{net}}$ is constant and the cost of persistence is strictly additive (section II E). The extension to flux-driven open systems is treated in section V.

B. The Three Primitives

Prigogine's Minimum Entropy Production theorem establishes *that* a persistent system must minimize its entropy production rate, but not *what* defines that rate. We derive the minimum dynamical description by identifying the irreducible primitives that any state transition law must presuppose.

When the static six-pillar architecture moves through time, any state transition (a local dissipative event) can be parameterized by three fundamental variables: what changes (a structured payload), what permits the change to occur (an energetic toll), and how often change occurs (a temporal frequency).

1. **A structured payload (The Subject):** A persistent system maintains a specific configuration rather than a statistical ensemble. Its local state is quantified by the structural complexity K: the Kolmogorov complexity of the local configuration relative to the substrate's reference encoding, expressed in substrate-natural units of elementary logical operations per local state (up to the invariant additive constant of the substrate's instruction set).[1] *Convergent Evidence (Kolmogorov Complexity [13]):* The algorithmic information content of a state is the irreducible payload carried by any transition.

2. **An irreducible toll (The Cost):** Every transition exacts a cost per elementary logical operation, c, expressed as a dimensionless number in fundamental units. *Convergent Evidence (Landauer's Principle [12]):* Every irreversible logical operation carries a thermodynamic cost of at least $k_B T \ln 2$.

3. **A coordinated rate (The Frequency):** State updates must occur at a transition flux f, a rate density of state updates per unit substrate measure. *Convergent Evidence (The Data-Rate Theorem [14]):* A dynamical system can be stabilized over a communication channel only if the information transmission rate exceeds the system's intrinsic entropy-production rate; below this threshold, the state estimate diverges and regulation is lost. f is the substrate-level analogue of this minimum rate, required to track the Map against environmental perturbation.

[1] Uncomputability does not preclude operational use: the algorithmic-information-theory literature routinely substitutes computable proxies: Minimum Description Length, Lempel-Ziv complexity, normalized compression distance, as finite-sample estimators of K, a practice the baseline derivation does not require but that applied instantiations can draw on directly.

The pillars operate on orthogonal domains (section II E) in the QSS; K, c, and f inherit this independence, each drawing on a distinct domain: algorithmic description, informational toll, and temporal rate, respectively.

C. The Dissipation Density

To compare competing configurations or define selection, the architecture requires a single scalar measure of dissipation. The three primitives must combine into such a measure. Dimensional bookkeeping constrains the form: a local dissipation metric must yield a uniform density of elementary transitions per unit substrate measure. These brackets track substrate counting units (operations, states, actions), not physical mass-length-time dimensions, the same sense in which a bit count or a radian is dimensionless yet still carried explicitly through a derivation. With K in [operations / local state], c in [entropic toll / operation], and f in [local states / (substrate measure · time)], their strictly linear product yields $\mathcal{D}$ in [entropic toll / (substrate measure · time)].

Dimensional bookkeeping alone does not fix the form uniquely. The dissipation metric must satisfy three joint requirements: vanishing when any primitive is absent; separability across orthogonal domains; and extensivity (under continuity) in each primitive independently.

Vanishing. $\mathcal{D}$ must vanish when any factor reaches zero, as shown in section III D: removing c, K, or f drives lifetime to zero. A non-vanishing form would permit dissipation without the corresponding pillar, contradicting the necessity proof that its absence causes collapse.

Separability. Because the three primitives draw on orthogonal domains (section II E), additive or non-separable couplings such as $c + K$ or e^{cK}, which mix dissipation across independent domains, are excluded. The functional form must be multiplicatively separable: $\mathcal{D} = g(c) \cdot h(K) \cdot k(f)$.

Extensivity. Doubling any one factor while holding the other two fixed must double $\mathcal{D}$. By Cauchy's functional equation, $g(\lambda c) = \lambda g(c)$ for all $\lambda > 0$ implies $g(c) = ac$ for some constant a, and likewise for h and k, ruling out both fractional-power forms such as $(c \cdot K \cdot f)^{1/3}$ and higher-power single-variable forms such as $cK^2 f$.

The general continuous, separable, extensive form is therefore $\mathcal{D} = A \cdot c \cdot K \cdot f$. The constant A is absorbed into the substrate's natural units (one elementary transition per unit substrate measure), yielding the canonical product:

$$\mathcal{D} = c \cdot K \cdot f \tag{2}$$

D. Necessity: Counterfactual Failure Modes

The necessity of (c, K, f) follows by the same counterfactual method used for the pillars (section II C): removing any factor drives lifetime to zero, as the local manifestation of its corresponding pillar failure mode.

- **No c (Zeno Collapse):** Without an irreducible cost per operation, the causal sequence loses its discrete granularity. The system decoheres into a continuum of indistinguishable micro-transitions. $\tau \to 0$, the local form of No-Toll.
- **No K (Dissolution):** No distinguishable local states remain; random activation maximizes uncertainty and the system–environment boundary dissolves. $\tau \to 0$, the local form of No-Map.
- **No f (Decoherence/Divergence):** The system freezes into a static configuration, decoupled from its own history and unable to mount a defensive response. $\tau \to 0$, the local form of No-Protocol/No-Governor.

Since the removal of any factor results in termination, the set (c, K, f) is necessary.

E. Sufficiency: No Fourth Factor

Given the multiplicative form, a missing factor would drive $\mathcal{D}$ to zero rather than merely reduce it, leaving no independent dimension for a fourth factor to occupy. Any property proposed as a fourth factor, either decomposes into combination of these three or specifies environmental boundary conditions rather than intrinsic system properties.

F. The Complexity Trade-off

The product form of the Dissipation Density implies a finite-budget competition among the three factors.

Proposition 1 (The Complexity Trade-off). *For any system in QSS, the substrate's finite Capacity imposes a dissipation ceiling $\mathcal{D}_{max}$:*

$$c \cdot K \cdot f \leq \mathcal{D}_{max} \tag{3}$$

No system can simultaneously maximize structural complexity, operational speed, and processing efficiency: increasing any one factor requires a proportional decrease in the product of the other two.

Proof. The substrate's finite Capacity imposes a hard ceiling $\mathcal{D}_{\max}$ on the total dissipation rate it can sustain. Exceeding this ceiling violates the bounded-resource condition of QSS, forcing the system out of steady state. Substituting the unique product form (equation (2)) into this ceiling yields the inequality directly. □

A biological cell illustrates the binding constraint: ATP cost per operation (c), proteome complexity (K), and metabolic turnover (f) together cannot exceed the thermodynamic ceiling of the cellular volume. A cell cannot maximize proteome complexity and reaction frequency at once without $\mathcal{D}$ diverging.

Convergent Evidence: As established for Capacity (section II B), the Margolus-Levitin limit and Bekenstein bound jointly cap the achievable volume in (c, K, f)-space, one bounding f for any finite c, the other bounding K.

G. The Entropic Action

Having established the conditions for structural survival, we need a way to compare surviving systems by lifetime. A system's viability cannot be judged pointwise: no single pillar sustains a transition alone, so no single local event determines survival either. Viability is the accumulation across the lifetime. We build that accumulated quantity in two steps: first collapsing the system's spatial extent into one number at each instant, then accumulating that number over its lifetime.

$\mathcal{D}$'s multiplicative, local form and $\bar{Z}_{\rm net}$'s additive, global one (section II E) are reconciled by integrating over the system's spatial footprint $\Sigma(t)$: many local multiplicative events, summed, become one global additive quantity at each instant:

$$Z_{\rm net}(t) = \int_{\Sigma(t)} \mathcal{D}\, d\mu \tag{4}$$

with dμ the counting measure on the discrete substrate state space, a sum over discrete points, not a continuous integral.

Since $Z_{\rm net}(t)$ is the total instantaneous dissipation, the cumulative **Entropic Action** is its time integral over the system's lifetime:

$$S_\Phi = \int_0^\tau Z_{\rm net}(t)\, dt \tag{5}$$

1. Physical Meaning of the Entropic Action

Because $\mathcal{D}$ counts elementary state transitions per unit substrate measure, S_Φ is a pure, dimensionless count of irreversible transitions, not yet an energy. When the dimensionless action S_Φ is projected onto a thermal physical substrate, Landauer's principle supplies the thermodynamic scale $c_0 = k_B T \ln 2$, yielding the physical dissipation analogue $E_{\rm diss} = c_0 S_\Phi$. On athermal substrates, S_Φ remains well-defined; only the conversion factor to macroscopic energy changes.

2. Necessity of the Architecture for the Action

Minimization of S_Φ is only possible for systems implementing the complete six-pillar architecture; the Entropic Action is not defined independently of it.

Proof. Persistence requires $\mathcal{D}$ strictly positive and finite, hence $c, K, f > 0$ (equation (2)). This requires all six pillars simultaneously: Toll and Capacity sustain $c > 0$; Map sustains $K > 0$; Protocol, Governor, and Margin jointly sustain $f > 0$. Removing any one pillar collapses the corresponding factor to zero (reproducing the failure modes of section III D), so any trajectory that maintains structural coherence necessarily encodes the complete architecture. □

The Entropic Action is therefore not defined independently of the architecture.

IV. MINIMUM ACTION AS SURVIVORSHIP

The Entropic Action defines what a system must minimize, but not which histories survive to be observed. We derive the selection filter from the multiplicative compounding of independent transition risk, and show that only minimal-action configurations remain observable.

A. The Exponential Filter

If persistence is maintained by state transitions, and each transition risks scrambling information into the substrate noise floor, then a history of N transitions survives only if every step survives. Independent errors compound multiplicatively, yielding an exponential filter.

Consider a structural history Γ accumulating entropic action $S_\Phi[\Gamma]$ (equation (5)). Because the substrate possesses a finite noise floor, each elementary transition carries an intrinsic risk of information loss. Let q be the probability that a single elementary transition maintains causal coherence.

In QSS, successive transitions are statistically independent dissipative events. The probability that Γ maintains coherence is therefore multiplicative:

$$P(\text{survival} \mid \Gamma) = q^{S_\Phi[\Gamma]} \tag{6}$$

Using the identity $q^C = e^{C \ln q}$ with $C = S_\Phi[\Gamma]$, and defining $\beta \equiv -\ln q > 0$ (positive since $0 < q < 1$), we have $S_\Phi[\Gamma] \ln q = -\beta S_\Phi[\Gamma]$, yielding the survival probability distribution:

$$P(\text{survival} \mid \Gamma) = e^{-\beta S_\Phi[\Gamma]} \tag{7}$$

Here $\beta \equiv -\ln q$ is the substrate's inverse entropic tolerance, a geometric invariant of the channel. For discrete substrates, q is determined by the channel geometry; in continuous media, β emerges from the substrate's noise spectral density.

This is the informational analogue of the Boltzmann distribution $P \propto e^{-E/k_B T}$, with S_Φ as energy and β as inverse temperature. The Boltzmann form requires ensemble equilibrium and a thermal bath; the present form requires only independent transition risk along a single history and a substrate noise floor. The informational route recovers the same exponential law under strictly weaker assumptions: substrate geometry replaces thermodynamic equilibrium.

B. The Severity of the Exponential Filter

The operational severity of this filter can be quantified by comparing trajectories. Let $\Gamma_{\min}$ denote the minimal-action trajectory and Γ_{var} any non-minimal path with an action deviation $\delta S_\Phi > 0$. Over any observation window, the relative survival probability is:

$$\frac{P(\Gamma_{\text{var}})}{P(\Gamma_{\min})} = e^{-\beta \delta S_\Phi} \tag{8}$$

When the fundamental update cycle is much shorter than the observation timescale, even a small per-transition deviation compounds into a vast accumulated δS_Φ, and the survival ratio $P(\Gamma_{\text{var}})/P(\Gamma_{\min})$ falls to zero. The gap between the best trajectory and the next-best widens without bound.

Because survival probability depends exponentially on S_Φ, only histories approaching the minimum retain non-negligible observation probability:

> The **Selection Principle**: Among all configurations capable of encoding a given structural complexity, only those possessing the minimal possible **Entropic Action** (S_Φ) remain observable.

Because S_Φ counts elementary transitions, it is the cumulative ledger of a system's dissipation history. The substrate does not optimize this ledger; it erases ledgers that fail to balance.

C. Eliminative vs. Teleological Selection

The severity of the exponential filter is a statement about the losers, not the winner. It says nothing about whether $\Gamma_{\min}$ is efficient, elegant, or well-suited to anything, only that every alternative to it decayed faster. A sub-optimal trajectory survives just as certainly as an optimal one, provided everything else was worse. What looks like design, foresight, or purpose is not a property of the winner; it is the absence of its competition.

D. Minimum Action as Survivorship

This eliminative mechanism is formalized as:

Proposition 2 (Minimum Action as Survivorship). *In QSS, the observable trajectories of any persistent system are those achieving the minimal Entropic Action S_Φ among all candidate trajectories.*

Proof. By equation (7), the survival probability of any history Γ is $P(\text{survival} \mid \Gamma) = e^{-\beta S_\Phi[\Gamma]}$. For any non-minimal trajectory Γ_{var} with $\delta S_\Phi = S_\Phi[\Gamma_{\text{var}}] - S_\Phi[\Gamma_{\text{min}}] > 0$, equation (8) gives $P(\Gamma_{\text{var}})/P(\Gamma_{\text{min}}) = e^{-\beta \delta S_\Phi}$. As δS_Φ accumulates over the observation window, this ratio vanishes. Therefore, only trajectories achieving the minimal Entropic Action retain non-negligible observation probability. The observable set is the argmin of $S_\Phi[\Gamma]$. □

Standard quantum mechanics grounds the Principle of Least Action in phase interference. The Selection Principle offers a distinct, classical mechanism (informational elimination) by which candidate physical histories are selected: the substrate itself acts as a filter, and what remains observable is only what has not yet dissipated.

V. OPEN SYSTEMS: NECESSITY WITHOUT ADDITIVITY

The Selection Principle and the Principle of Least Action were derived in the Quiescent Steady State: the zero-flux limit where no net energy or information enters or leaves. Many systems, however, exchange flux with their environment to persist. We show that the necessity of the six pillars survives this extension unchanged; only the net impedance (Z_{net}, section II E) requires re-derivation in flux-driven capacity.

A. Pointwise Necessity and Operational Latency

Because the substrate continuously filters for causal coherence, survival is an instantaneous requirement, not a time-averaged one.

Proposition 3 (Pointwise Necessity). *For any system persisting past instant t, all six pillars must be structurally instantiated at t, regardless of external flux.*

Proof. By the counterfactual failure modes of section II C and section III D, the absence of any pillar drives the corresponding primitive to zero ($c \to 0$, $K \to 0$, or $f \to 0$), causing the system's expected lifetime to vanish ($\tau \to 0$). For any finite interval $\mathrm{d}t$, the probability of surviving from t to $t + \mathrm{d}t$ therefore tends to zero. The requirement is instantaneous and flux-independent. □

Flux alters a pillar's *operation*, not its *existence*. During phases of rapid growth or unbounded resource influx, a control mechanism may exhibit low activation ($f_{\text{GOV}} \approx 0$), minimizing active coordination expenditure. Yet its *structural capacity* must exist at t; if the mechanism is fundamentally absent, the system fails to persist the moment it reaches environmental saturation.

B. Open Systems and the Dynamic Budget

Strict additivity ($\bar{Z}_{\text{net}} = \sum_i Z_i$) relies on the Quiescent Steady State, where orthogonal domains prevent interference. In open systems, external flux raises the dissipation ceiling; as operations scale against this moving ceiling, resource contention forces non-linear cross-coupling $\Delta(t)$ between previously orthogonal pillars:

$$Z_{\text{net}}(t) = \sum_i Z_i(t) + \Delta(t). \tag{9}$$

A stochastic derivation of $\Delta(t)$ from substrate-native properties is developed in appendix A; three regimes emerge as its limits, where ρ is the dimensionless coordination load:

1. **Linear-Response Regime (Low ρ):** $\Delta \to 0$; QSS additivity is the leading-order behavior.
2. **Saturation-Divergence Regime (High ρ):** Coordination cost diverges as the dissipation ceiling is approached, forcing structural pauses (apoptosis, circuit breakers) to shed load.
3. **Correlated-Demand Regime (Non-Markovian):** Relaxing the memoryless assumption recovers the standard heavy-traffic generalization, with the $\rho/(1-\rho)$ form as its memoryless special case.

The derivation does not alter the necessity claim. The six pillars remain structurally mandatory regardless of the specific form of $\Delta(t)$; the idealized limit where $\Delta(t) \to 0$ recovers QSS additivity.

That derivation also provides a unified account of these three regimes, which were previously analyzed independently. Linear-response recovery, heavy-traffic divergence, and the correlated-demand generalization were each discovered independently in unrelated fields. Here they emerge as three limits of one substrate function.

VI. UNIVERSALITY MAPPING: ARCHITECTURE AND DYNAMICS

The core mechanics of persistence are identifiable across systems of any substrate. Mapping both the static architecture (the six pillars) and the operational dynamics (the dissipation trade-off and selection filter) translates domain-specific phenomena into identical mechanical constraints. The following mappings demonstrate architectural universality across distinct physical and informational domains. Detailed operationalizations belong to applied persistence mechanics; here, we illustrate the underlying structural invariants.

A. Biological Cell

Formalizing the biological concept of autopoiesis [15]:

Architecture: Capacity as metabolic hardware; Map as genome; Protocol as genetic code and signaling machinery; Governor as apoptosis and cell-cycle checkpoints; Toll as the thermodynamic cost of proofreading; Margin as ATP reserves and stress buffers.

Dynamics: The cell's basal metabolic rate ($\mathcal{D}$) is the product of its ATP cost per molecular reaction (c), proteome complexity (K), and metabolic turnover rate (f). An organism cannot simultaneously maximize genomic complexity and reproductive speed without exceeding the local thermodynamic ceiling. Trajectories that overdraw this budget undergo exponential suppression via starvation or competitive exclusion.

Cross-Domain Transfer: Rather than treating aging as a correlational pattern of stochastic damage, this architectural perspective reframes senescence as the progressive degradation of operational pillars across nested sub-systems, revealing structural failure modes analogous to resource exhaustion in computational systems.

B. Ecological Network

Formalizing ecological ascendency and network thermodynamics [16]:

Architecture: Capacity as biomass and solar energy capture; Map as trophic structure and niche topology; Protocol as nutrient cycling and metabolic exchange; Governor as carrying-capacity and predator-prey feedbacks; Toll as thermodynamic heat loss per trophic transfer; Margin as biodiversity and structural redundancy.

Dynamics: Total ecosystem respiration ($\mathcal{D}$) is the product of heat lost per transfer (c), food-web complexity (K), and the rate of biomass turnover (f). Because solar influx rigidly caps $\mathcal{D}$, overly complex food webs with inefficient trophic transfers are systematically pruned by famine and ecological collapse, leaving only the least-action trophic topology.

Cross-Domain Transfer: Mapping ecological networks to energy grids and financial markets reveals that cascading collapse obeys the same fixed dissipation ceiling, making grid load-shedding and market circuit-breakers candidate tools for anticipating tipping points in biodiversity loss.

C. Large Language Model

Formalizing a young domain still establishing its foundations:

Architecture: Capacity as inference compute and GPU memory bandwidth; Map as parametric weights and attention matrices; Protocol as tokenization, layer connections, and key-value (KV) caching; Governor as output constraints and behavioral guardrails; Toll as autoregressive latency and state-update friction; Margin as unused context window, processing redundancy, and decision margin.

Dynamics: The model's active dissipation burden ($\mathcal{D}$) is the product of compute friction per generated token (c), context and parameter footprint (K), and generation velocity (f). In autoregressive generation, attempts to maximize context length and reasoning depth without lowering parameter activation costs hit a hard memory bandwidth and compute ceiling. The Entropic Action acts as a strict evolutionary filter on AI architectures: models that fail to minimize this product succumb to cumulative attention noise (hallucination) or catastrophic inference costs, driving the structural selection of minimal-action survivors (e.g., KV-cache distillation, sparse routing, or linear-time state-space models).

Cross-Domain Transfer: The comparatively young field of AI can draw on cross-disciplinary knowledge without re-deriving it: long-context degradation is currently benchmarked and treated empirically. If the Saturation-Divergence Regime correctly describes KV-cache saturation, its onset is a hazard-rate problem already solved by Weibull analysis.

Contention Regimes: The three regimes of appendix A are visible in inference. Short-context generation resembles the Linear-Response Regime. As context length approaches the memory-bandwidth ceiling, generation quality collapses in a manner resembling the Saturation-Divergence Regime, the KV-cache saturation and attention degradation already documented in long-context models. Non-uniform, correlated attention patterns (repetition, recency bias) degrade worse than a memoryless baseline would predict, suggesting the Correlated-Demand Regime as a candidate lens, not a demonstrated one.

D. Further Applications

The same six-pillar mapping extends to domains further removed from physical substrates. Two examples illustrate distinct payoffs of that extension, briefly rather than in full: one shows the mapping makes previously incomparable systems commensurable; the other shows the architecture binds even a domain with no physical substrate at all.

- **Institutional Organizations:** A shared dissipation metric shows how bureaucratic decay in a growing firm and signal decay in a fiber-optic network are directly comparable structural failures.

- **Scientific Enterprise:** Reframing Planck's observation that science advances one funeral at a time as the thermodynamic survival of funded believers over verification friction shows that theoretical paradigms are constrained by the same dissipation limits as cells or firms.

VII. CONCLUSION

Fluctuation theorems, master equations, and Langevin dynamics describe the statistics of transitions; the architecture of persistence describes the structure that must exist before any such statistics can be defined, the discrete foundation that any continuous dynamics must presuppose. For the fields of adaptation and self-organization, this shifts the fundamental question: before a system can adapt or self-organize, it must possess a minimal architecture permitted to exist.

We have derived that architecture and the mechanism that selects for it. The six pillars are necessary; the Entropic Action is the filter; and the Selection Principle emerges as the cumulative boundary condition governing surviving configurations.

This mechanics supplies a formal boundary condition for persistence that complex systems research has treated implicitly but not defined: the minimal architecture any adaptive or self-organizing system must instantiate before adaptation or self-organization can occur. Because this universal mapping projects every persistent system onto the same six-pillar substrate, it functions as a Rosetta stone across disciplines. Hard-won solutions in one domain become transferable theorems in another: protocol engineering that prevents signal decay in a fiber-optic network also prevents fragmentation in a corporate hierarchy. In emerging domains such as AI, where the pillars are present but unnamed, this vocabulary lets systems be engineered for resilience without re-deriving structural knowledge from scratch.

The same mapping enables direct quantitative comparison among competing systems in a shared environment. By evaluating how competing organisms, firms, or technical architectures manage their finite dissipation budgets, relative operational resilience can be assessed prior to empirical failure.

Organizational decline, ecosystem collapse, and the limits of biological longevity have been described through domain-specific frameworks. By defining exactly what is required to persist, this mechanics recasts them as structural failures under a single, unified mechanics: disparate phenomena become comparable instances of the same dissipation constraint, rather than loose analogies.

The Complexity Trade-off is a pre-observational constraint on structural complexity, coordination speed, and energetic efficiency. A cell, a firm, or a model pushing two of these to its limits has already revealed where the third must be constrained; any two fix the third before measurement begins.

The most consequential limiting case is fundamental physics: whether standard physical law is itself a substrate-specific instantiation of this same invariant. Further exploration will determine whether this architectural boundary holds universally. Candidate mathematical languages for testing this include Information Geometry (continuous substrates), Arithmetic Geometry (discrete substrates), and Sheaf Theory (compositional substrates).

Natural avenues for the extension include systems composed of nested persistent subsystems, what it takes for a system to arise from disorder, and a fuller account of decay itself: this paper treats termination as $\tau \to 0$, without characterizing the trajectory a system follows on the way there, or whether decay has its own structure worth deriving.

Persistence Mechanics does not explain why structure arises. It explains why only certain architectures remain to be observed. Persistence requires no intent, no optimization, and no forward-looking agency, only that the substrate continuously filter non-viable configurations. What remains observable is the minimal architecture of persistence.

ACKNOWLEDGMENTS

The author is an independent researcher and received no external funding for this work. I would like to thank James Hughes and Brian Sheppard for rigorous and constructive feedback. I would like to thank my friends and family for their patience and support throughout decades of discussions, as I tried to understand every field I became interested in and the recurring, overarching patterns I saw.

[1] E. Schrödinger and R. Penrose, *What is Life?: With Mind and Matter and Autobiographical Sketches*, 1st ed. (Cambridge University Press, 1992).

[2] I. Prigogine, *Introduction to Thermodynamics of Irreversible Processes* (Interscience Publishers, New York, 1961).

[3] K. Friston, The free-energy principle: a unified brain theory?, Nature Reviews Neuroscience **11**, 127 (2010).

[4] J. L. England, Dissipative adaptation in driven self-assembly, Nature Nanotechnology **10**, 919 (2015).

[5] W. R. Ashby, *An Introduction to Cybernetics* (Chapman & Hall, London, 1956).

[6] R. C. Conant and W. Ross Ashby, Every good regulator of a system must be a model of that system, International Journal of Systems Science **1**, 89 (1970).
[7] J. D. Bekenstein, Universal upper bound on the entropy-to-energy ratio for bounded systems, Physical Review D **23**, 287 (1981).
[8] N. Margolus and L. B. Levitin, The maximum speed of dynamical evolution, Physica D: Nonlinear Phenomena **120**, 188 (1998).
[9] H. Nyquist, Regeneration Theory, Bell System Technical Journal **11**, 126 (1932).
[10] J. F. C. Kingman, The single server queue in heavy traffic, Mathematical Proceedings of the Cambridge Philosophical Society **57**, 902 (1961).
[11] C. Shannon, Communication in the Presence of Noise, Proceedings of the IRE **37**, 10 (1949).
[12] R. Landauer, Irreversibility and Heat Generation in the Computing Process, IBM Journal of Research and Development **5**, 183 (1961).
[13] A. N. Kolmogorov, Three approaches to the quantitative definition of information, International Journal of Computer Mathematics **2**, 157 (1968).
[14] G. N. Nair, F. Fagnani, S. Zampieri, and R. J. Evans, Feedback Control Under Data Rate Constraints: An Overview, Proceedings of the IEEE **95**, 108 (2007).
[15] H. R. Maturana and F. J. Varela, *Autopoiesis and Cognition: The Realization of the Living*, Boston Studies in the Philosophy and History of Science, Vol. 42 (Springer Netherlands, Dordrecht, 1980).
[16] R. E. Ulanowicz, *Growth and Development* (Springer New York, New York, NY, 1986).

Appendix A: The Dynamic Budget: Contention for Shared Resources

Open systems violate the orthogonality that made $\bar{Z}_{\rm net}$ strictly additive in QSS. To quantify the resulting contention between Protocol and Governor, we derive the cross-coupling $\Delta(t)$ in equation (9) from discrete substrate properties.

1. Coordination capacity and load

Protocol and Governor share the substrate that performs all state transitions. The Margolus-Levitin limit bounds the minimum time between distinguishable operations, yielding a substrate-level coordination rate $f_{\max}$. In QSS, the two pillars operate on orthogonal domains (their demands share no microstate), so their loads add. With $f_{\rm PRO} + f_{\rm GOV} \ll f_{\max}$, cycles remain effectively empty and orthogonality holds ($\Delta = 0$). The dimensionless coordination load is therefore

$$\rho(t) = \frac{f_{\rm PRO}(t) + f_{\rm GOV}(t)}{f_{\max}} \tag{A1}$$

2. Temporal independence and the geometric waiting distribution

The load above governs instantaneous rates, not waiting-time statistics. We model retry behavior as a memoryless process: each cycle presents pending events symmetrically, with service probability $1-\rho$ per event. This is a modeling hypothesis about temporal independence, not a consequence of orthogonality. The waiting time follows a geometric distribution with mean

$$\langle k \rangle = \sum_{k=0}^{\infty} k\,\rho^k (1-\rho) = \frac{\rho}{1-\rho}. \tag{A2}$$

3. The coordination cross-term

Each pending cycle incurs the fixed toll c_0 (the dimensionless Landauer scale; on thermal substrates this projects to $k_B T \ln 2$). The total waiting cost for both streams is

$$\Delta_{\rm total}(t) = \frac{c_0\, f_{\max}\, \rho}{1-\rho}. \tag{A3}$$

The QSS baseline (equation (1)) already prices the linear coordination costs $Z_{\rm PRO} + Z_{\rm GOV} \approx c_0(f_{\rm PRO} + f_{\rm GOV})$ at $\rho \ll 1$. The cross-coupling $\Delta(t)$ is the excess beyond this baseline:

$$\Delta(t) = \Delta_{\rm total}(t) - c_0\, f_{\max}\, \rho = \frac{c_0\, f_{\max}\, \rho^2}{1-\rho}. \tag{A4}$$

This recovers the three regimes of section V B: linear-response recovery as $\rho \to 0$, saturation divergence as $\rho \to 1$, and correlated-demand behavior under heavy-tailed waiting distributions.

4. Scope and validity

This construction assumes discrete cycles (the Margolus-Levitin bound) and symmetric, memoryless retry. Strongly correlated, heavy-tailed, or long-range-dependent arrivals fall outside its scope.